Distribution of elliptic twins over fixed finite fields: Numerical results

David Leon Gil

September 11, 2015

1 Preface

This is a preliminary note of some numerical experiments; the results may be rather wrong.

2 Introduction

This paper presents the results of numerical experiments to determine the probability, over concrete fixed finite fields, of prime-order elliptic curves having a prime-order twist.

These curves are called "elliptic twins" by [7], and are useful for a variety of cryptographic applications.[1]

Most notable is that such curves are secure against an "insecure twist" attack. This attack was introduced in 2001 by Daniel Bernstein, see [2], who has proposed "twist-security" (a slightly weaker condition) as an essential safety criterion for elliptic curves. [1][2]

The most interesting result of this paper is that, for the finite fields the NSA-generated curves are defined over, there is only an approximately 1/100 probability of a random prime-order curve having a prime-order twist.

P-384 was standardized by NIST in 1999, and generated by the NSA at some previous time.[9] It has a prime-order twist. [2]

P-224 was standardized by NIST at the same time. It does not have a prime-order twist. In fact, its twist has only 58-bit security. [3]

[1] [7] only consider the asymptotic density of "elliptic twins" as a fraction of all elliptic curves, so their results only partially address the question of this paper. One analytic approach might be to combine their results with the results of [**GalbraithMcKee**].

[2] [1] autocites Burton Kaliski ([5]) as introducing the so-called "unsafe-twist" attack, but I have been unable to find any evidence either in that paper or Kaliski's thesis, [4], that he was aware of the attack. Kaliski's construction of an elliptic-curve-and-twist-based random number generator does, however, require that discrete log be hard on both the curve and its twist, as he explicitly notes.

[3] The twist of P-224 has a cofactor of $3^2 \cdot 11 \cdot 47 \cdot 3015283 \cdot 40375823 \cdot 267983539294927$. [2]

3 Elliptic twin curves

We follow the definitions of [7], with some minor modifications.

Let $\mathbb{E}_j$ be the elliptic curve of invariant j, and $\mathbb{E}_j(\mathbb{F}_q)$ be its reduction over a finite field of characteristic $p > 5, n \geq 1$ with p prime. Let and $t(\mathbb{E}_j(\mathbb{F}_q))$ be the trace of Frobenius of that elliptic curve.

Let $\widetilde{\mathbb{E}}_j(\mathbb{F}_q)$ be the non-trivial quadratic twist of $\mathbb{E}_j(\mathbb{F}_q)$ over the same field.

An *elliptic twin* is a pair consisting of a prime p, and a set of two primes not equal to p or 0, $\{l, r\}$, such that

$$\#\mathbb{E}_j(\mathbb{F}_p) + \#\widetilde{\mathbb{E}}_j(\mathbb{F}_p) = l + r = 2p + 2 - l + r \tag{1}$$

[7] provide evidence that elliptic twins exist over arbitrary prime fields, but the formulae of [7] do not appear to provide precise estimates for fixed finite fields.

4 Primes

We consider the non-Mersenne SECP primes, standardized for the use of the federal government in [9], which are, where $N := 2^{32}$:

$$\begin{aligned} \mathrm{P}_{224} &= N^7 - N^3 + N^0 \\ \mathrm{P}_{256} &= N^8 - N^7 + N^6 + N^3 - N^0 \\ \mathrm{P}_{384} &= N^{12} - N^4 - N^3 + N^1 - N^0 \end{aligned} \tag{2}$$

They are subset of the class of Generalized Mersennes defined by [8].

5 Numerical methods

5.1 Finding prime-order curves

Method 1. A slightly modified version of PARI/GP was used to calculate the traces of prime-order curves, based on code of [**HamburgPARI**]. (The particular code used for this version of this paper may be found at [6].) Point-counting was aborted early if $\#\mathbb{E}_j$ was found to have a small prime factor. This computation produced estimates for the density both of prime-order curves and elliptic twin curves over each field.

Method 2. For P-384, a slightly larger computation was carried out using the same code, but set up to abort point-counting if either $\#\mathbb{E}_j or \#\mathbb{E}^t_j(\mathbb{Z}_q))$ had a small prime factor. The experiment

5.2 Results

Experiment 1. We calculate $T_f(j)$ for each $\mathbb{E}_j$ for $0 < j < 2^{20}, j \neq 1728$, then test $\#\mathbb{E}_j$ and $\#\mathbb{E}^t_j(\mathbb{Z}_q))$ for primality.

For this to be a reasonable procedure, it requires the assumption that j-invariant is not correlated with the probability of the curve being an elliptic twin, even on a local scale of 2^{20}.

Let N_π be the number of prime-order curves found, and $N_{\pi'}$ be the number of elliptic twins found. Then, in this range, we have:

	N_π	$N_{\pi'}$	$N_{\pi'}/N_\pi$
P_{224}	2790	31	1.1e-2
P_{256}	1956	15	0.8e-2
P_{384}	1131	20	1.8e-2

Experiment 2. Because of the small number of elliptic twin curves found in expriment 1, we planned to conduct the following experiment: For 1000 pseudorandomly-generated j-invariants, set $j_{i,0}$ to a j-invariant, and increment until $j_{i,n}$ is an elliptic twin. The average of $j_{i,n}$ is then an estimator for $1/p(pi')$.

Due to resource constraints this experiment was aborted after finding only 441 elliptic twins.

Combining these results with those of experiment 1, after bootstrapping, give a 99% confidence interval for $p(pi'|pi) = [0.005, 0.01]$.

6 Future work

In future work, we plan to extend the study to consider the more general question of the distribution of group structure and curve exponent for reductions of curves over fields for which their number of integral points is non-prime, and apply similar techniques with respect to the two curves proposed for IETF use, the nearly-Mersenne $\mathrm{M}_{255} = 2^{255} - 19$ and the Hamburg-Solinas trinomial $\mathrm{H}_{448} = 2^{448} - 2^{224} - 1$. [4]

(We probably won't extend this work to the Mersenne M_{521}, as that particular calculation is pestiferously large.)

7 Conclusion

The quantity $1/\mathrm{p}(p) = N_\pi(\mathrm{p})/N_{\pi'}(\mathrm{p})$ is an estimator for the number of trials required, when choosing a prime curve uniformly at random in $\mathbb{F}_q$ for that curve to be an elliptic twin.

The probability, however, that no elliptic curves in a set of N curves are elliptic twins is, of course,

$$1 - \left(\prod_{0 \leq i < n} (1 - \mathrm{p}_i)\right) \tag{3}$$

[4] The Hamburg primes are "Karatsuba-friendly" and [3] was the first to publish an algorithm that fully takes advantage of their special form.

With respect to the curves generated by the NSA for [**SECP1**], and subsequently standardized by [9], this calculation gives a probability of very approximately $> 95\%$ that *none* of the curves over $\mathrm{P}_{224}, \mathrm{P}_{256}$, and P_{384} would be an elliptic twin.

But the curve over P-384 *is* an elliptic twin.

One might thus conclude that it is more likely than not that the NSA's curves were not generated by a process that samples from a uniform distribution on prime-order curves over the chosen prime fields.[5]

In particular, this suggests that the NSA's choice of seeds for the "random" prime curves were subject to additional safety criteria not yet publicly disclosed. (Or, of course, that things with 5% probability aren't terribly rare events...[6])

In addition, it suggests that the fever for "twist-security" which has taken grip of the cryptographic community is potentially dangerous: These are a smallish class of elliptic curves, and there is no evidence that – provided an implementation is not vulnerable to a small-twist attack – they possess either more or *less* structure than a non-twist-secure curve.

8 Acknowledgments

This work was inspired by Daniel Bernstein's SafeCurves website, and a frustratingly long search for a twist-secure curve over M_{607}.

Many thanks to Robert Ransom for his skeptical comments, which have helped clarify the argument of this note.

The patch to PARI/GP is derived from a patch by Michael Hamburg.

Appendix. Cofactors for SafeCurves

This table is adapted (read stolen directly) from [2]. The rows have been sorted by the cofactor of the twist of the curve. The curves for which twist-security was a stated security criterion during the selection process have been omitted.

Curve	$h(\mathbb{E}_j)$	$h(\mathbb{E}_j^t(\mathbb{Z}_q)))$
`secp384r1`	1	1
`secp256r1`	1	$3 \cdot 5 \cdot 13 \cdot 179$
`secp256k1`	1	$3^2 \cdot 13^2 \cdot 3319 \cdot 22639$
`FRP256v1`	1	$7 \cdot 439 \cdot 11760675247 \cdot 3617872258517821$
`secp224r1`	1	$3^2 \cdot 11 \cdot 47 \cdot 3015283 \cdot 40375823 \cdot 267983539294927$
`brainpoolP256`	1	$5^2 \cdot 175939 \cdot 492167257 \cdot 8062915307 \cdot 2590895598527 \cdot 4233394996199$
`brainpoolP384`	1	$7 \cdot 11^2 \cdot 241 \cdot 5557 \cdot 125972502705620325124785968921221517$

[5] Why, then, don't all of P_{224} , P_{256} , and P_{384} have safe twists? Note that the probability of that would be $\prod_{0 \leq i < n}(1 - \mathrm{p}_i)$, or less than $1.5e - 6$, or a roughly 1 in 630,000 chance.

[6] In partial defense of the NSA: Suppose that it did, in fact, draw the seeds for the SECP prime curves uniformly at random until it found prime order curves. There is no good way of the NSA "proving" that it followed this procedure honestly, even if it did. This reinforces the importance of some "rigidity" criterion, as per [**NUMS**].

References

[1] Daniel J Bernstein. “Curve25519: new Diffie-Hellman speed records”. In: *Public Key Cryptography-PKC 2006*. Springer, 2006, pp. 207–228.

[2] Daniel J. Bernstein. *SafeCurves: choosing safe curves for elliptic-curve cryptography: Twist security*. Oct. 2013. URL: http://safecurves.cr.yp.to/twist.html.

[3] Mike Hamburg. “Fast and compact elliptic-curve cryptography”. In: (2012). URL: http://eprint.iacr.org/2012/309.pdf.

[4] Burton Stephen Kaliski. “Elliptic curves and cryptography: A pseudorandom bit generator and other tools”. PhD thesis. Massachusetts Institute of Technology, 1988.

[5] Burton S Kaliski Jr. “One-way permutations on elliptic curves”. In: *Journal of Cryptology* 3.3 (1991), pp. 187–199.

[6] Michael Hamburg PARI Authors and David Leon Gil. *coruus/junk-pari*. URL: https://github.com/coruus/junk-pari.

[7] Igor E Shparlinski and Daniel Sutantyo. “Distribution of elliptic twin primes in isogeny and isomorphism classes”. In: *Journal of Number Theory* 137 (2014), pp. 1–15.

[8] Jerome A. Solinas. *Generalized mersenne numbers*. Tech. rep. 99-39. Center for Applied Cryptographic Research, University of Waterloo, 1999. URL: http://cacr.uwaterloo.ca/techreports/1999/corr99-34.pdf.

[9] National Institute of Standards and Technology. *Recommended elliptic curves for federal government use*. 1999.